KubeCon CloudNativeCon India 2026

#KubeCon #CloudNativeCon

Service Networking Within Air-Gapped Environments: Deployment Strategies and Operational Management

Anirban Nandi, Google

Agenda

01 What is an air-gapped environment aka a “private cloud”?

02 Typical application requirements in such platforms

03 How does Istio help?

04 Deployment model

05 Operational Challenges

The Air-Gapped Platform

Google Distributed Cloud: Air-Gapped

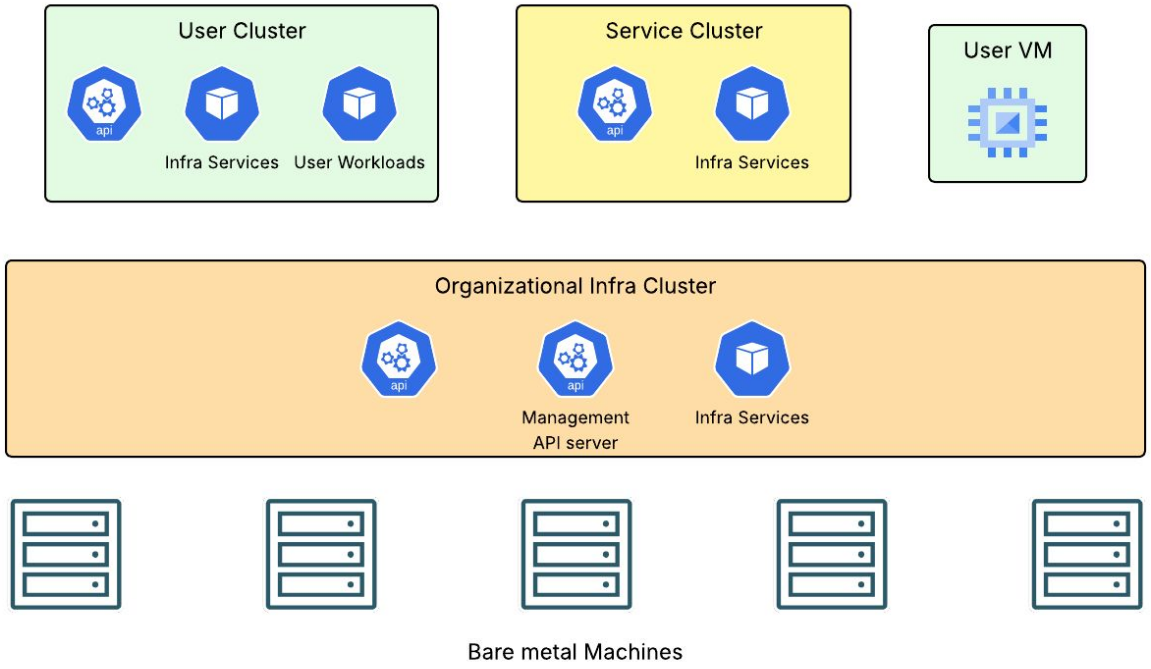

- HW+SW solution delivering modern cloud capabilities without connectivity to public cloud
- Offerings include Virtual Machines Services, IAM, Key Management, Database Services, Storage, Networking, etc.
- K8s based orchestration with OSS solutions for infrastructure services.
- Organization and Project like tenancy with different personas for infrastructure, platform and application operations.

KubeCon + CloudNativeCon North America 2024: How Google Built a New Cloud on Top of K...

Core Application Requirements

Observability

- Metrics like QPS, latency, throughput, etc.
- Audit logs and tracing for traffic across multiple hops.

Security

- Encryption of data in transit
- Authorized access only
- DOS prevention (e.g. rate limiting)

Traffic Management

- L7 attribute based routing
- Custom load balancing across backends
- N-S and E-W reachability
- Traffic mutation (e.g. header transformation, gRPC/HTTP transcoding, etc.)

Imagine having to do all of these for **45+ services each having 10 microservices** on an average

Service Networking using Istio

What is Istio?

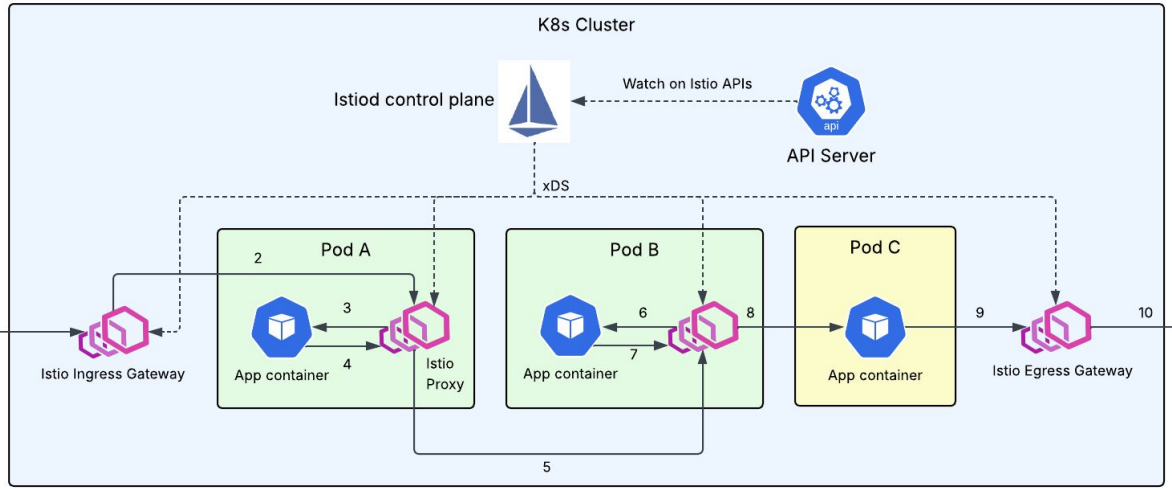

- Service Mesh with **Envoy** as dataplane
- Two modes: **Sidecar** (in picture) and **Ambient**
- Configuration propagation using **xDS**
- **VirtualService, DestinationRule, Gateway**, etc. APIs to configure policies
- Supports multi-cluster, multi-network, multi-primary topologies
- Offers traffic shifting, mirroring, circuit-breaking, auto mTLS, telemetry, etc.

Istio in GDC Air-Gapped

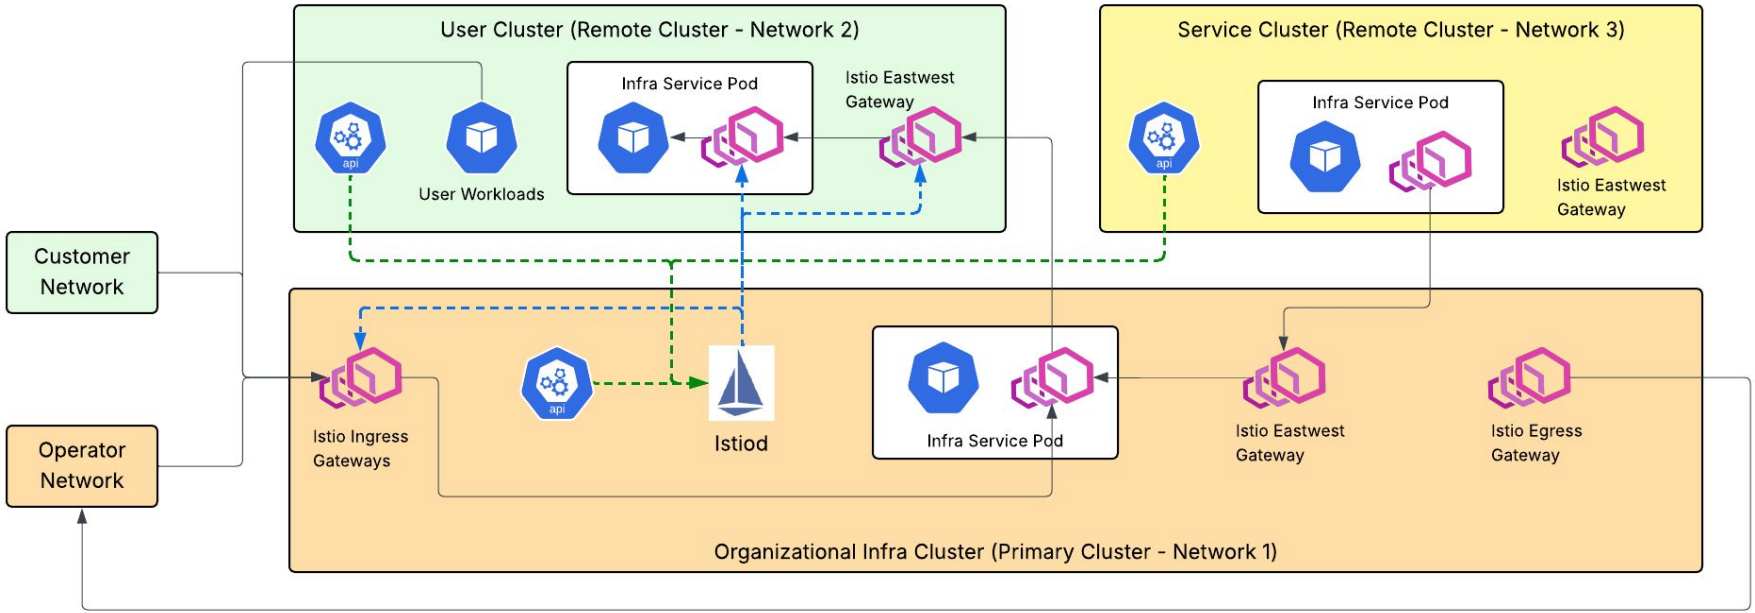

Gateways

Edge & Transit Proxy

- North-South: Customer to Infra Services
- East-West: Between Infra services
- Expose API servers securely

Authentication

- TLS termination (**Gateway API**)
- TLS origination (**DestinationRule API**)

Authorization

JWT, CIDR-based, and external authorizers

(**RequestAuthentication, AuthorizationPolicy API**)

Load Balancing

- L7 based routing (**VirtualService API**)
- Custom LB (**DestinationRule API**)

Observability

Application metrics: QPS, Throughput, Latency, etc.

(**Telemetry API**)

Application Security

Mutual TLS

Automatic mTLS for secure service-to-service communication.

(PeerAuthentication API)

Certificate Management

Management with plug-in CA for flexible identity control.

(IstioOperator API)

Policy Control

Granular Allow/Deny policies for traffic governance.

(AuthorizationPolicy API)

Application Telemetry

OTel Audit Logs

Standardized auditing and logging for requests traversing the mesh.

(Telemetry API)

Prometheus Metrics

Rich Istio and Envoy stats for tracking performance and defining SLIs.

(Telemetry API)

Selective Metrics

Flexible inclusion and exclusion of metrics at the pod level.

(Pod annotations)

Application Traffic Management

L7 Routing

Advanced HTTP/gRPC routing based on headers, paths, and methods.

(VirtualService API)

Traffic Shifting

Gradual migration from old to new versions via weighted load balancing.

(VirtualService API)

Circuit Breaking

Outlier detection and fail-fast mechanisms to ensure system stability.

(DestinationRule API)

Fault Injection

Testing resiliency by injecting delays or aborting requests intentionally.

(VirtualService API)

Dataplane Extensibility

Use Envoy features even without first-class Istio APIs (EnvoyFilter API)

Lua
Lightweight header and payload transformations.

Wasm
Complex scripts in sandboxed VMs for high performance.

Transcoding
Native HTTP/gRPC transcoding support.

Rate Limiting
Apply global and local rate limiting policies.

Compression
Gzip, Brotli, and more.

External Auth
Delegate auth decisions.

… many more filters available in Envoy

Lifecycle Management

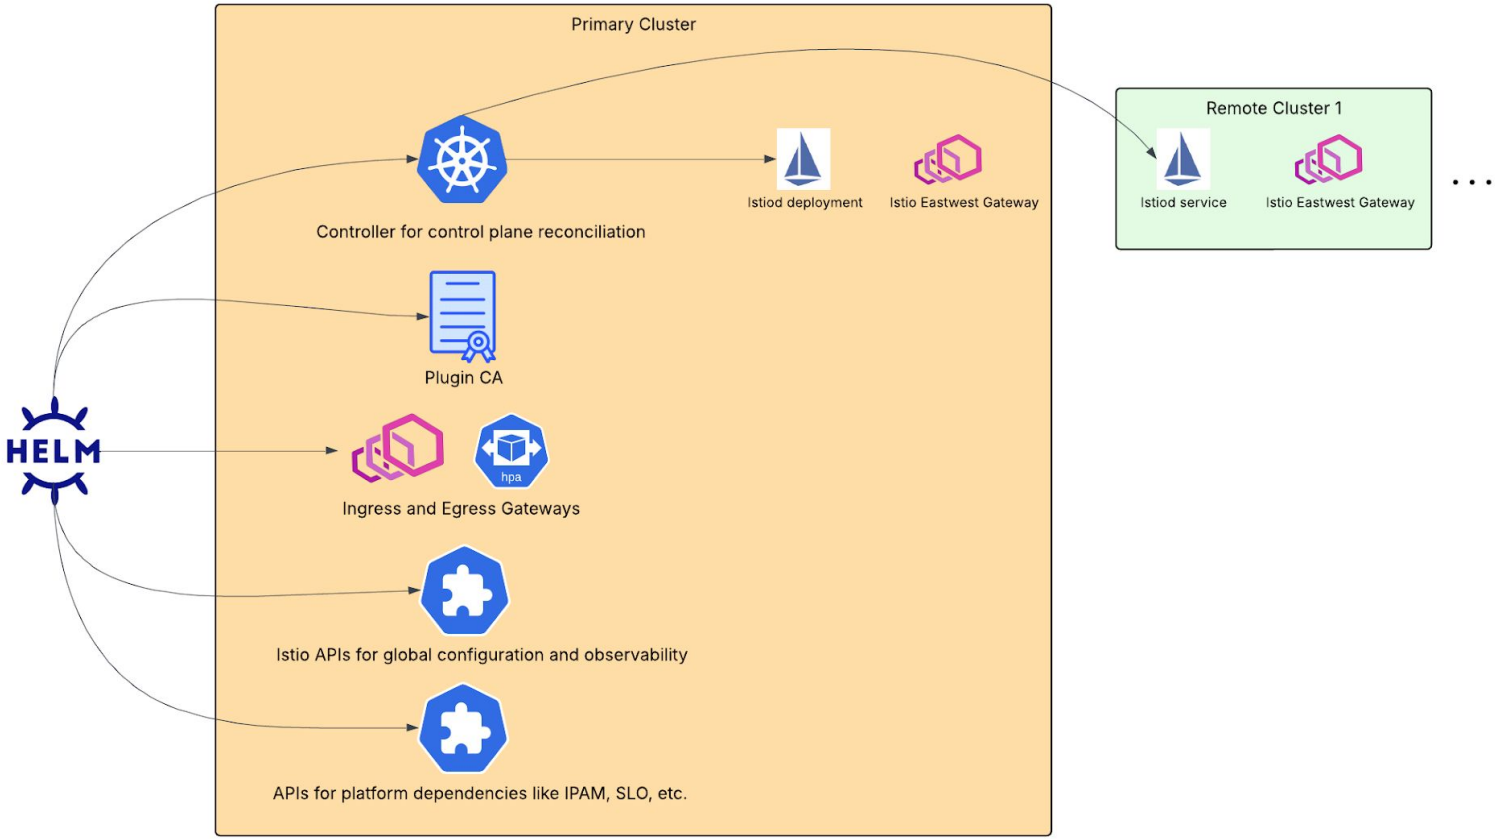

Integration with OSS tools

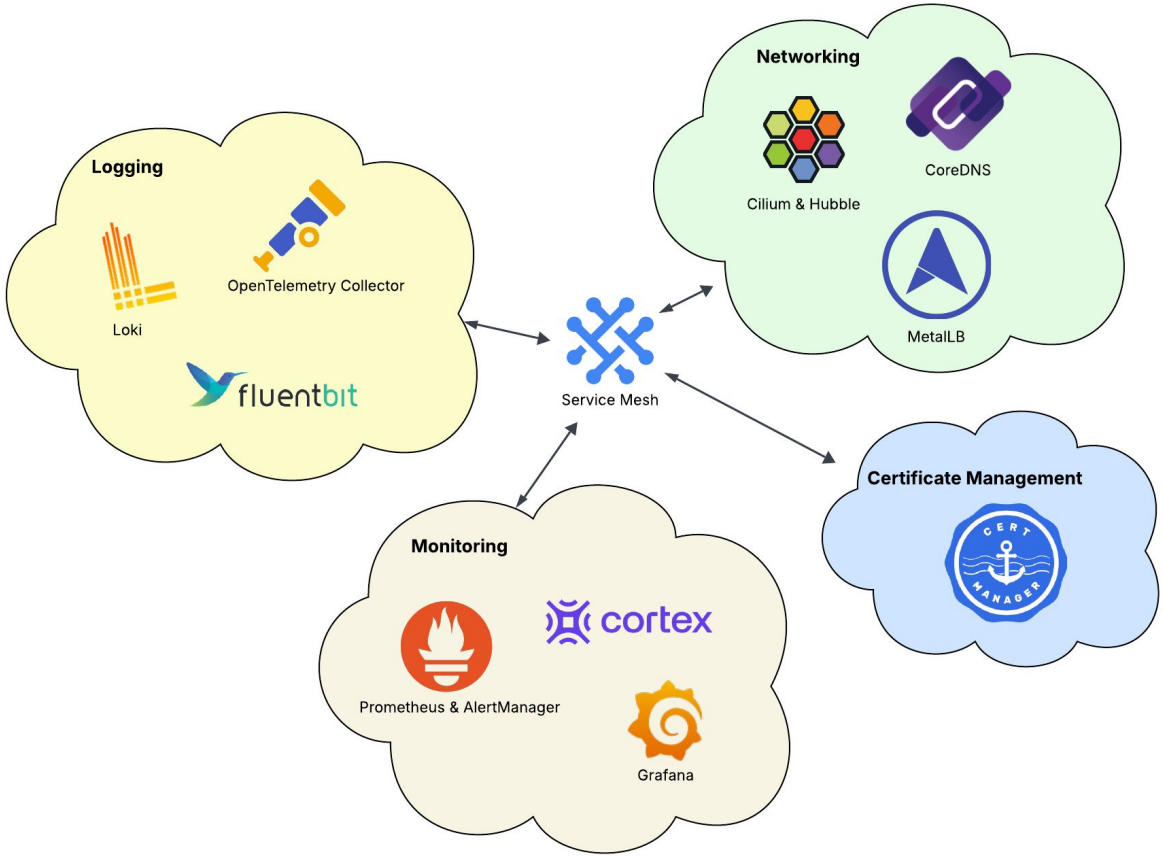

- Cilium for L4 networking
- MetalLB for external IP allocation and advertisement
- CoreDNS for DNS forwarding across clusters
- Loki, Fluentbit, OTel Collector for capturing mesh audit logs and operational logs.
- Grafana for viewing metrics and operational logs
- Prometheus and Cortex for capturing Istio and Envoy stats
- AlertManager for registering alerts
- Cert Manager for plugin CA

Operational Challenges

Advanced Feature set

- No support via first class APIs
- Direct xDS modification conflicts with Istiod generated xDS

Debuggability

- Selective metric aggregation to prevent cardinality explosion
- Selective audit and proxy log aggregation to prevent storage overload

Co-existence with 3P workloads

- Resource conflicts with 3P Istio deployments
- Prevent intra-mesh communication between 1P and 3P workloads

Scalability

- Large config propagation contained by config scoping and delta xDS
- Constrained by the underlying bare metal compute

Lifecycle Management

- Ordered upgrade of mesh components in primary-remote multi-cluster
- Special handling for unhealthy user clusters
- Minimal downtime for dataplane during upgrades
- Auto-scaling of gateways

QnA